

Europäisches Patentamt
European Patent Office
Office européen des brevets



(11) EP 0 713 183 A2

(12)

EUROPEAN PATENT APPLICATION

(43) Date of publication:
22.05.1996 Bulletin 1996/21

(51) Int. Cl.⁶: G06F 17/30

(21) Application number: 95117949.8

(22) Date of filing: 14.11.1995

(84) Designated Contracting States:
DE FR GB

(30) Priority: 18.11.1994 US 342127

(71) Applicant: MICROSOFT CORPORATION
Redmond, Washington 98052-6399 (US)


(72) Inventor: Pardikar, Shishir
Redmond, Washington 98052 (US)

(74) Representative: Grünecker, Kinkeldey,
Stockmair & Schwanhäusser
Anwaltssozietät
Maximilianstrasse 58
80538 München (DE)

(54) Network independent file shadowing

(57) Network independent file shadowing is provided by the present invention. The file shadowing mechanism automatically and transparently stores shadow copies of remote file system structures when they are accessed by a computer. The shadow copies of the file system structures are stored within a shadow database that

resides within local memory of the computer. When the computer becomes disconnected from a network, shadow copies of file system structures for the network are used to service requests to access such file system structures.



BEST AVAILABLE COPY

EP 0 713 183 A2

Description

Technical Field

The present invention relates generally to data processing systems and, more particularly, to network independent file shadowing.

Background of the Invention

One difficulty encountered with portable computers concerns accessing files when disconnected from a network. Most portable computers provide facilities for the portable computer to be connected to a network. When the portable computer is connected to the network it may gain access to remote files that are stored within the network. Unfortunately, when the portable computer becomes disconnected from the network, the portable computer no longer has access to the remote files stored in the network.

Summary of the Invention

In accordance with a first aspect of the present invention, a method of providing access to a file system structure from a network when a computer is disconnected from the network is provided. The method is practiced in a data processing system that includes a computer with a storage device. The computer is connectable to a first network of a first type and a second network of a second type. In this method, when the computer is connected to a select of one of the networks and a program is running on the computer, several steps are performed. First, a request to access one of the selected file system structures stored in the selected network is received from the program. In response to this request, a shadow copy of the selected file system structure is transparently obtained so that the shadow copies are obtained transparently relative to the program. The shadow copy of a selected file system structure is stored in a shadow database on the storage device. The shadow database holds file system structures for use when the computer is no longer connected to the selected network.

In accordance with another aspect of the present invention, a method is practiced in a data processing system that has a computer with a storage. The computer is connectable to a plurality of networks of different types. At least one of the networks stores a selected file system structure. In this method, a shadow database is provided in the storage for storing shadow copies of file system structures that are retrieved for the networks. The shadow database includes a shadow copy of the selected file system structure from the selected network. When the computer is disconnected from the selected network, a request to access the selected file system structure is received and the shadow copy of the selected file system structure that is stored in the shadow database is used to service the request.

In accordance with a further aspect of the present invention, a method is practiced in a data processing system having a computer with a storage and the computer is connectable to a plurality of networks of different types. Each network includes at least one file system structure. In this method, when the computer is connected to the network, requests to access file structures of the networks are received and shadow copies of the file system structures for which requests to access are received are stored in a shadow database in the storage. When the computer is disconnected from at least one of the networks and a request to access a selected file system structure in one of the disconnected networks is received, the shadow database is accessed to use a shadow copy of the selected file system structure to service the request.

In accordance with yet another aspect of the present invention, a computer system is connectable to different types of networks. The computer system includes a shadow database that holds shadow copies of selected file system structures from different types of networks. The computer system also includes a shadow module for responding to a request or access one of the selected file system structures by accessing the shadow database to obtain a shadow copy of the requested file system structure to service the request. The computer system may also include a shadow database builder for building the shadow database. The shadow database builder stores shadow copies of selected file system structures when the file system structures are accessed by the computer. The computer system also may include an agent for building and maintaining the shadow database.

Brief Description of the Drawings

Figure 1 is a block diagram of a distributed system that is suitable for practicing the preferred embodiment of the present invention.

Figure 2 is a block diagram illustrating the software components that play a role in practicing the preferred embodiment of the present invention.

Figure 3 is a flowchart providing an overview of the steps performed in the preferred embodiment of the present invention.

Figure 4 is a flowchart illustrating the steps that are performed to have the shadow VxD hook file input/output requests in accordance with the preferred embodiment of the present invention.

Figure 5 is a flowchart illustrating the steps that are performed when a file I/O request is made for a remote file in accordance with the preferred embodiment of the present invention.

Figure 6 is a flowchart illustrating steps that are performed by the shadow database in accordance with the preferred embodiment of the present invention.

Figure 7 is a flowchart illustrating the steps that are performed by the agent to remove stale entries in the

shadow database in the preferred embodiment of the present invention.

Figure 8 is a flowchart illustrating the steps performed by the agent to create sufficient free space in the shadow database in accordance with the preferred embodiment of the present invention.

Figure 9 is a flowchart illustrating the steps that are performed to use the shadow database when a computer is disconnected from a network in accordance with the preferred embodiment of the present invention.

Detailed Description of the Invention

The preferred embodiment of the present invention provides network independent file shadowing. The file shadowing provided by the preferred embodiment is network independent in that files from different types of networks may be shadowed using the mechanism of the preferred embodiment of the present invention. The file shadowing performed by the preferred embodiment of the present invention automatically and transparently keeps copies of files that are accessed by a user in a remote network file system. The transparency provided by the preferred embodiment relieves the user of the burden of selecting which files to store locally for access when the computer is disconnected from the remote networks.

Figure 1 is a block diagram of an illustrative distributed system 10 that is suitable for practicing the preferred embodiment of the present invention. Those skilled in the art will appreciate that the distributed system depicted within Figure 1 is intended to be merely illustrative and that the present invention may be practiced in other data processing system configurations. The distributed system 10 depicted in Figure 1 includes networks 12, 14, and 16. Each of these networks includes a file server 18, 20 and 22, respectively. The file servers 18, 20 and 22 regulate access to files and other data stored within the respective networks 12, 14, and 16. Each of the networks 12, 14, and 16 may include additional components such as workstations, storage devices, and other peripheral devices. Network 12 includes a workstation 24 with memory 26. The memory 26 holds a copy of an operating system 28 and at least one application program 30. Code for practicing the preferred embodiment of the present invention is encapsulated within the operating system 28. Nevertheless, those skilled in the art will appreciate that in alternative embodiments this code may be encapsulated into separate modules that are not directly incorporated into the operating system.

For purposes of the discussion below, it is assumed that the workstation 24 may be decoupled from networks 14 and 16. Moreover, it is assumed that the workstation 24 is a portable computer that is readily adaptable for use in an isolated environment. Those skilled in the art will appreciate that the present invention is not limited to practice on a portable computer, but rather may also be practiced on other types of computer systems.

The preferred embodiment of the present invention overcomes the difficulties of the prior art described within the Background section by automatically and transparently storing shadow copies of files that are accessed by the workstation 24 from remote networks (like networks 14 and 16). As will be described in more detail below, the preferred embodiment of the present invention provides a file shadowing mechanism and a database of remotely accessed files that serve as a persistent cache. The file shadowing mechanism hooks requests for file input/output (I/O) from the remote networks. The files involved in such requests are then brought into the persistent cache. Thus, while workstation 24 is connected to networks 14 and 16, a persistent cache is built so that the files within the persistent cache may later be accessed when the workstation is no longer connected to networks 14 and 16.

Figure 2 is a block diagram illustrating software components that play a role within the preferred embodiment of the present invention. Border 37 separates the software components at the user privilege level from software components at the kernel privilege level. The file I/O requests are generated by the application program 30. The application program 30 may be any type of application, including a word processing or spreadsheet program. Those skilled in the art will appreciate that the present invention also is applicable to instances wherein the file I/O requests are initiated from within the operating system 28, rather than from within an application program 30. For purposes of the discussion below, an application program is intended to include any program that can generate a file I/O request.

The request originating from the application program 30 is passed to a multiple provider router 31. The multiple provider router 31 routes the request to the network providers 33 and 35. The network providers 33 and 35 are provided as client code in the operating system 28. The network providers 33 and 35 include functions that allow a connection to be created between the application program 30 and a remote network. The multiple provider router 31 passes the requests to each of the network providers to determine which network provider can successfully connect with the appropriate network. Separate network providers 33 and 35 are provided for each of the different types of networks within the distributed system 10. The preferred embodiment of the present invention provides a shadow network provider which is used when none of the other network providers 33 succeed in establishing a connection. The shadow network provider 35 is used to establish a shadow network connection to the shadow database 42 so that the data contained therein may be accessed. The use of the shadow network provider 35 will be described in more detail below. The network providers 33 and 35 may establish a connection to a remote network by directly interacting with the file system drivers 38 and 40, which will be described in more detail below. Alternatively, the network providers 33 may establish connections to

remote networks via the installable file system manager 32.

The IFSMGR 32 is included as part of the operating system 28 and is responsible for directing file I/O requests to their proper destination (*i.e.*, to the proper file server for handling the request). The IFSMGR 32 is designed to work with multiple file systems so that a user may access files in different file system types.

The IFSMGR 32 interacts with file system drivers (FSDs), like FSDs 36, 38, and 40, which are shown in Figure 2. The FSDs 36, 38, and 40 serve as interfaces between the IFSMGR 32 and the file systems. The FSDs appear to the system as device drivers. Specifically, the FSDs 36, 38, and 40 are viewed as device drivers that are responsible for a file system.

The FSDs 36, 38, and 40 are part of the operating system 28. The FSDs 36, 38, and 40 may interface with local file systems (like FSD 36) or remote file systems (like FSDs 38 and 40). Each of the FSDs 36, 38, and 40 is designed to interface with a given file system.

The shadow VxD 34 is a virtual device driver for monitoring file I/O requests to determine what files need to be placed within the persistent cache (known as the shadow database 42). The shadow VxD 34 is a virtual device driver that simulates a device (*i.e.*, it represents the virtualization of a device). The shadow VxD 34 is implemented as part of the kernel of the operating system 28. The responsibilities of the shadow VxD 34 will be described in more detail below.

A user hints database may be provided to determine what file type should be included and excluded from the shadow database 42. The hints database acts like a filter to include or exclude certain file references. The hints database also acts like a work queue for the agent in that the agent may periodically fill in the requested file references in the hints database, regardless of whether the files have been referenced by any other application program. A hints database tries to maximize the utility of the local disk space in order to optimize the space that is available for the user when the user's computer is disconnected from a network.

The shadow database 42 holds complete naming information for files that are accessed by the workstation 28 (Figure 1) and that are within a remote file system. The shadow database 42 also holds a representation of the file that is accessed (*i.e.*, either a sparse representation or a complete representation of the file). The shadow database 42 also has a priority queue wherein higher priority files are at the head of the queue and lower priority files are at the end of the queue. Higher priority files are those which have a higher priority for remaining within the shadow database 42 than lower priority files.

An application, known as an agent 44, is provided to maintain the shadow database 42. As will be described in more detail below, the agent 44 is periodically awoken to bring files into the shadow database 42 and to update the shadow database as needed.

Figure 3 is a flowchart showing an overview of the steps that are performed by the preferred embodiment

of the present invention. Initially, a workstation 24 is connected to networks 14 and 16 (step 46 in Figure 3). The shadow database 42 is then built as will be described in more detail below (step 48 in Figure 3). Later, the workstation 24 becomes disconnected from one or more of the networks 14 and 16 (step 50 in Figure 3). While the workstation 24 remains disconnected from the networks, the shadow database 42 is used as a persistent cache that is accessed to locate files that originated from remote file systems (step 52 in Figure 3). The workstation 24 then reconnects to the networks (step 54 in Figure 3) and various clean-up activities are performed (step 56).

The shadow VxD 34 is responsible for hooking the file I/O requests to remote file systems. Figure 4 is a flowchart that illustrates the steps that are performed to enable the shadow VxD 34 to perform such hooking. Initially, during system boot, the IFSMGR 32 is loaded into memory 26 of the workstation 24 (step 58 in Figure 4). Once the IFSMGR 32 is fully loaded, the shadow VxD 34 is loaded (step 60 in Figure 4). The shadow VxD 34 hooks the calls to the RegisterNet service, which registers the remote file system drivers with the IFSMGR 32 (step 62 in Figure 4). The remote file system drivers 38 and 40 are then loaded subsequently (step 64 in Figure 4). The RegisterNet service registers remote networks with the IFSMGR 32. By hooking this call to the RegisterNet service, the shadow VxD 34 can hook each file I/O request that is sent from the IFSMGR 32 to one of the remote file system drivers 38 and 40.

When each of the file system drivers 38 and 40 are loaded, they call the RegisterNet service that is provided by the IFSMGR 32 to register the FSDs. The FSDs 38 and 40 forward a function pointer to make a connection to the file system driver as part of a call to the RegisterNet service. Subsequently, when the IFSMGR 32 needs a connection to be made, it calls the function that is identified by the pointer that is passed from the FSD 38 or 40. When this function call succeeds, the FSD returns an array of function pointers to the IFSMGR 32. This array of function pointers points to functions that can be performed on the file system associated with the FSD 38 or 40. The shadow VxD 34 hooks the RegisterNet call from the FSD 38 and 40 and receives the array of function pointers. The shadow VxD 34 passes its own array of function pointers to the IFSMGR 32. The array of function pointers passed to the IFSMGR from the shadow VxD 34 provide hooks into functions of the shadow VxD 34. As a result, each time that an operation is to be performed on one of the remote file systems, the shadow VxD 34 hooks these calls and fills the shadow database 42 accordingly.

Figure 5 is a flowchart illustrating the steps that are performed by the preferred embodiment of the present invention when an I/O request originates from application program 30. The process is initiated by the application program 30 making a file I/O request to a file that resides in one of the remote networks (step 66 in Figure 5). As was discussed above, this remote file I/O request

may originate from any program that is running on the local workstation 24. The request is passed to the IFS-MGR 32 (step 68 in Figure 5). The shadow VxD 34 hooks the request (step 70 in Figure 5) and adds the file to the shadow database 42 if it is not already present within the shadow database (step 70 in Figure 5). In particular, complete naming information for the requested file is stored within the shadow database 42. This complete naming information includes a pathname for the server-share pair as specified according to the unified naming convention (UNC). The files are stored within the database 42 in a sparse format (*i.e.*, the contents of the file are not stored immediately therein). The shadow VxD 32 then passes the file I/O request to the proper FSD so that the appropriate file is retrieved (step 74).

As was mentioned above, the role of the agent 44 (Figure 2) is to maintain the shadow database 42. One of the primary roles of the agent 44 is to retrieve the contents of files that are stored sparsely within the shadow database 42. The agent 44 performs this task as a background thread. Figure 6 is a flowchart illustrating the steps that are performed by the agent 44 in this capacity. Initially, the agent 44 is awoken on a periodic basis (step 76 in Figure 6). For example, the agent 44 may be awoken every 30 seconds. It is then active and subsequently put back to sleep. When the agent is awoken it looks at the entries within the shadow database 42 and identifies entries that are stored in sparse form (step 78 in Figure 6). The agent 44 begins looking for sparse entries at the top of the priority queue and then proceeds down to lower priority entries. When the agent 44 locates a sparsely stored entry, the agent attempts to fill in the contents of the file by grabbing a shadow copy of the file (step 80). The agent 44 is then put back to sleep after the fixed period of time elapses or after the agent has nothing further to do (step 82).

The agent 44 is also responsible for maintaining the shadow database 42. One difficulty is that copies of files stored within the shadow database 42 may become stale or too old. Each time a file is retrieved and stored within the shadow database 42, a timestamp is entered in the database. This timestamp is used as shown in Figure 7 to remove stale entries from the shadow database 42. In particular, the agent 44 examines the timestamp for files stored within the shadow database 42 (step 84 in Figure 7). The shadow VxD may perform this examination on opening of the files in the shadow database or at other suitable times. The timestamp is compared to a predetermined timestamp value to determine if the file is stale or not (step 86). Any timestamps with values earlier than the cutoff timestamp value are targeted for removal. If a file is determined to be stale, a new copy of the file is retrieved and stored within the shadow database 42 (step 88 in Figure 7).

Another difficulty that the agent 44 must address in maintaining the shadow database 42 is that the size of the shadow database 42 is fixed and hence, may become full. For example, the shadow database 42 may constitute a fixed size portion of a hard disk drive for the

workstation 24. As more files are brought into the shadow database, there is an increasing likelihood that the shadow database 42 will become full. In such an instance, certain files must be removed from the shadow database. As shown in Figure 8, the agent 44 determines that insufficient space is available (step 90 in Figure 8). This determination may be made upon writing to the shadow database 42, upon each opening of a file in the shadow database, or at other alternative times. The agent 44 begins deleting the lowest priority files on a one-by-one basis until sufficient space in the shadow database 42 is available (step 92 in Figure 8).

As was discussed above, the shadow database 42 is built while the workstation 24 is connected to the networks 16 and 18. The shadow database 42 is used, however, primarily when the workstation 24 is disconnected from one or more of the networks 16 and 18. Figure 9 is a flowchart illustrating the steps that are performed to use the shadow database 42 when the workstation 24 is disconnected from one or more of the networks 16 and 18. The application program 30 is running on the workstation 24 and initiates a file I/O request for a file on a remote file system (step 94 in Figure 9). The multiple provider router 31 determines that the conventional network providers 33 cannot be used to realize a connection to the network that holds the file in the remote file system and, thus, routes the request to the shadow network provider 35, which calls the shadow VxD 34 to access the shadow database 42 (step 96 in Figure 9). The shadow VxD 34 then becomes active and looks within the shadow database 42 (step 98 in Figure 9). The shadow VxD 34 checks whether the file is in the shadow database 42 and present (*i.e.*, not in sparse form) (step 100 in Figure 9). If the file is stored within the shadow database 42, it is retrieved and used (step 102 in Figure 9). On the other hand, if the file is not within the database or is merely stored in sparse form, an error indicating that the file is not available is returned (step 104).

When the workstation 24 is reconnected with one or more of the networks 14 and 16, a file server 18 and 20 of the network must examine the local copies of files held at the workstation 24 and reconcile these copies with those held in the network. Standard reconciliation techniques may be utilized. The agent 44 may include code for checking whether a given file server is online or not. When the server is online and local files have been modified, the integration and reconciliation is triggered.

While the present invention has been described with reference to a preferred embodiment thereof, those skilled in the art will appreciate that various changes in form and detail may be made without departing from the intended scope of the present invention as defined in the appended claims.

Claims

1. In a data processing system having a computer with a storage device and that is connectable to a first network of a first type and a second network of a

second type, wherein the first network and the second network each store file system structures, a method of providing access to one of the file system structures from one of the networks when the computer is disconnected from at least one of the networks, comprising the steps of:

when the computer is connected to a selected one of the networks and a program is running on the computer,

- (i) receiving a request to access a selected one of the file system structures stored on the selected network from the program;
- (ii) in response to the request, transparently obtaining a shadow copy of the selected file system structure relative to the program; and
- (iii) storing the shadow copy of the selected file system structure in a shadow database on the storage device which holds file system structures connected to the selected network.

2. The method of claim 1, further comprising the steps of:

when the computer is disconnected from the selected network,

- (i) receiving a request to access the selected file system structure of the selected network; and
- (ii) using the shadow copy of the selected file system structure that is stored in the shadow database to service the request.

3. The method of claim 2, further comprising the step of flushing the shadow database to remove all file system structures for the selected network when the computer is reconnected to the selected network.

4. The method of claim 1, further comprising the steps of:
- before storing the shadow copy, determining that there is insufficient free space in the shadow database to store the shadow copy; and
 - deleting file system structures in the shadow database to create sufficient free space to store the shadow copy of the selected file system structure.

5. The method of claim 1 wherein the shadow database is a priority queue that stores file system structures in the priority queue according to priority of the file system structures and wherein the selected file system structure has a priority and said step of storing the shadow copy of the selected file system structure at a location in the priority queue is based on the priority of the selected file system structure.

6. In a data processing system having a computer with a storage and being connectable to a plurality of networks of different types wherein at least a selected one of the networks stores a selected file system

structure, a method of providing access to the selected file system structure when the computer is disconnected from the selected network, comprising the steps of:

- providing a shadow database in the storage for storing shadow copies of file system structures that are retrieved from the networks, said shadow database including a shadow copy of the selected file system structure from the selected network;
- when the computer is disconnected from the selected network,

- (i) receiving a request to access the selected file system structure; and
- (ii) using the shadow copy of the selected file system structure that is stored in the shadow database to service the request.

7. The method of claim 6 wherein the step of providing the shadow database comprises the step of providing the shadow database as a priority queue wherein each shadow copy of a file system structure has a priority and a position of each shadow copy within the queue is determined by its priority.

8. The method of claim 6 wherein the computer system runs an application program and wherein the application program generates the request to access the selected file system structure and wherein said step of using the shadow copy of the selected file system structure to service the request is performed transparently relative to the application program.

9. The method of claim 6, further comprising the steps of:

when the computer is disconnected from a second of the networks,

- (i) receiving a request to access a file system structure in the second network; and
- (ii) using a given shadow copy of the file system structure to be accessed from the second network to service the request, the given shadow copy being stored in the shadow database.


10. In a data processing system having a computer with a storage and being connectable to a plurality of networks of different types wherein each network includes at least one file system structure, a method comprising the steps of:

when the computer is connected to the networks,

- (i) receiving requests to access file system structures in the networks;
- (ii) storing shadow copies of the file system structures for which requests to access are received in a shadow database in the storage;

when the computer is disconnected from at least one of the networks and a request to access a selected file system structure in one of the disconnected networks is received, accessing the shadow database to use a shadow copy of the selected file system structure stored therein to service the request. 5

11. The method of claim 10 wherein each shadow copy of a file system structure stored in the shadow database has a priority and said step of storing shadow copies of the file system structures for which requests to access are received in the shadow database in the storage comprises the step of storing the shadow copies as a priority queue wherein a position of each shadow copy within the priority queue is determined by its priority. 10
15
12. The method of claim 10 wherein the step of storing shadow copies comprises the steps of: 20
determining that insufficient free space is available in the shadow database to store a given one of the shadow copies;
deleting shadow copies from the shadow database to provide sufficient free space to store the given shadow copy. 25
13. The method of claim 10, further comprising the steps of:
storing a timestamp for each shadow copy in the storage, each timestamp indicating when the shadow copy was stored in the shadow database; 30
storing updated versions of shadow copies when the timestamps indicate that the presently stored shadow copies are too old. 35
14. A computer system that is connectable to different types of networks comprising:
a shadow database holding shadow copies of selected file system structures from different types of networks; and 40
a shadow module for responding to a request to access one of the selected file system structures by accessing the shadow database to obtain a shadow copy of the one file system structure to service the request. 45
15. The computer system of claim 14, further comprising a shadow database builder for building the shadow database by storing shadow copies of the selected file system structures when the file system structures are accessed by the computer. 50
16. The computer system of claim 14, further comprising an agent for maintaining the shadow database. 55
17. The computer system of claim 14, further comprising file system drivers for interfacing with file systems of the networks.



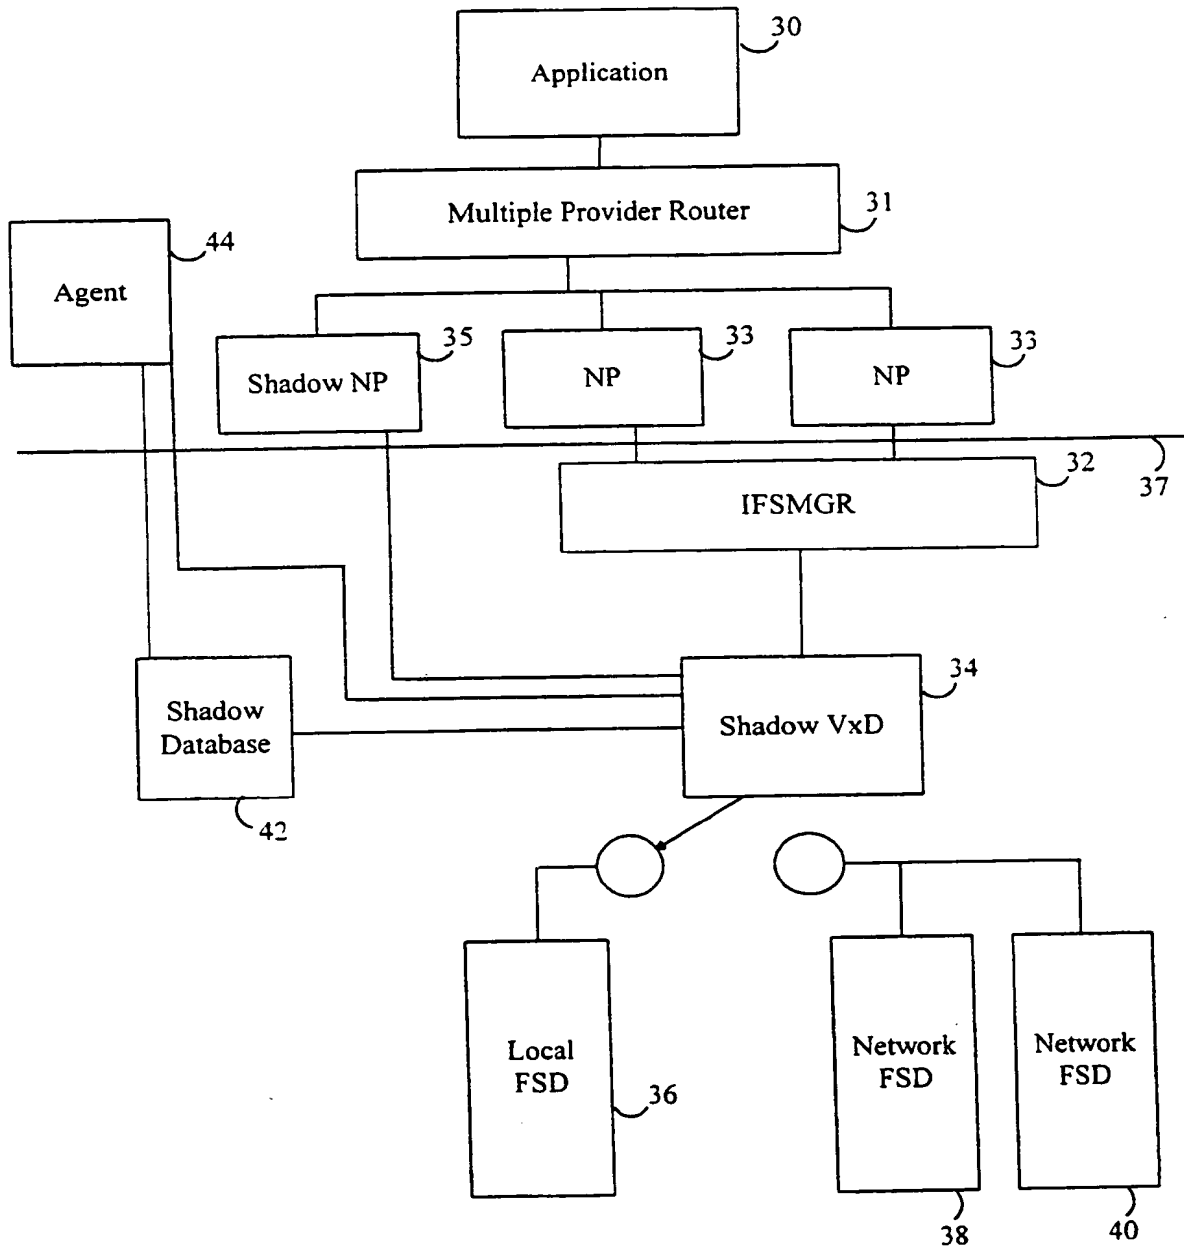


FIG. 2

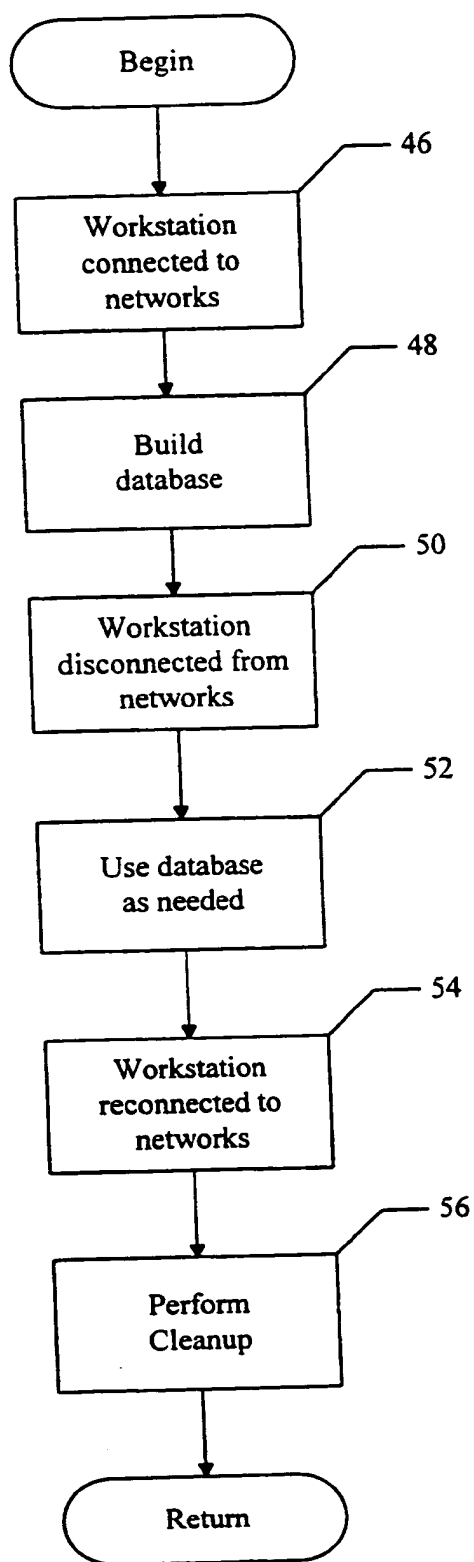


FIG. 3

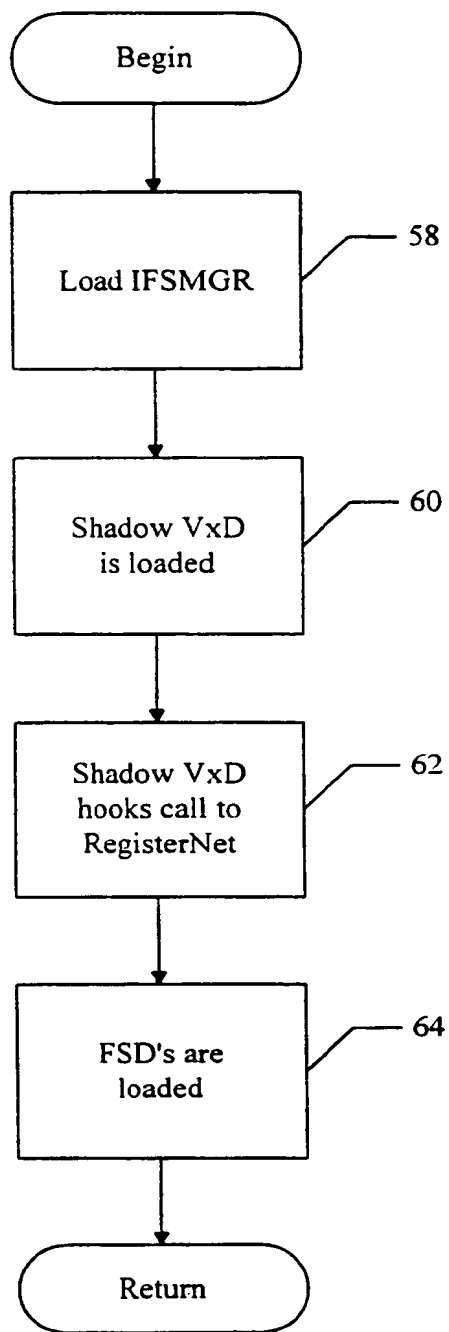


FIG. 4

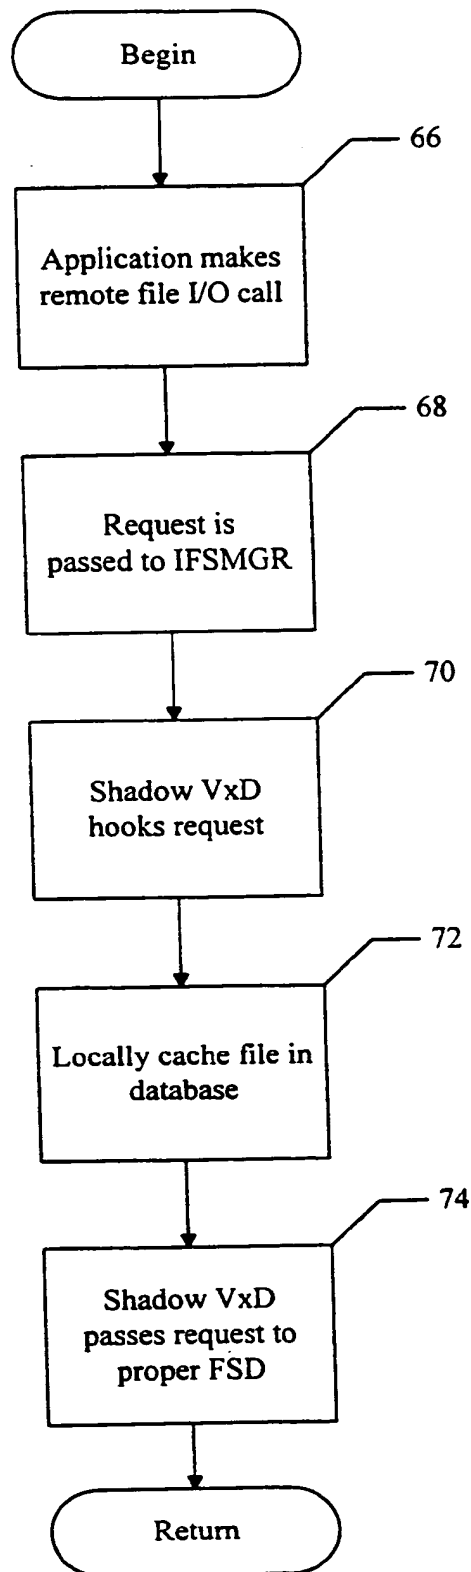


FIG. 5

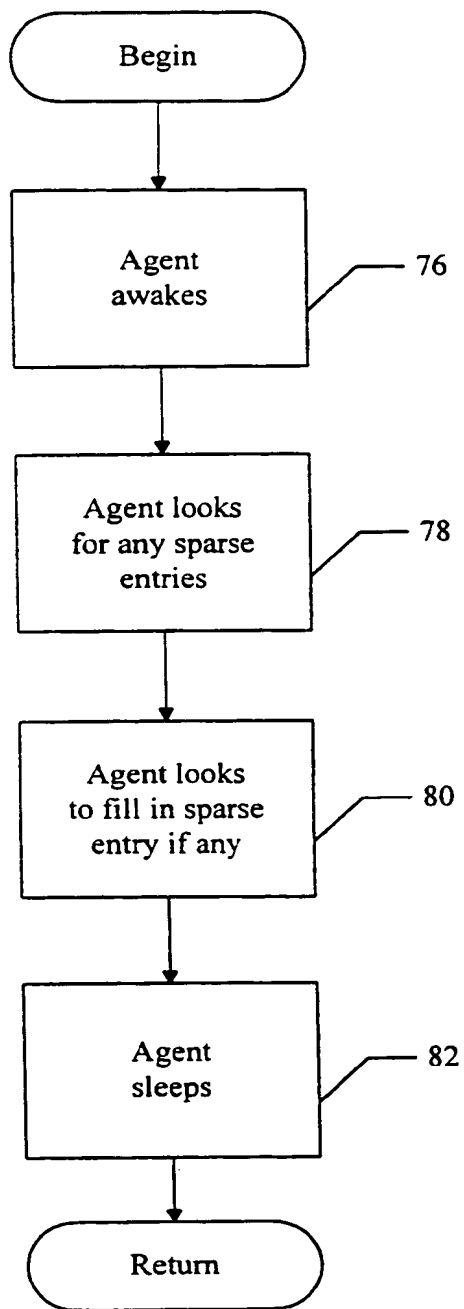


FIG. 6

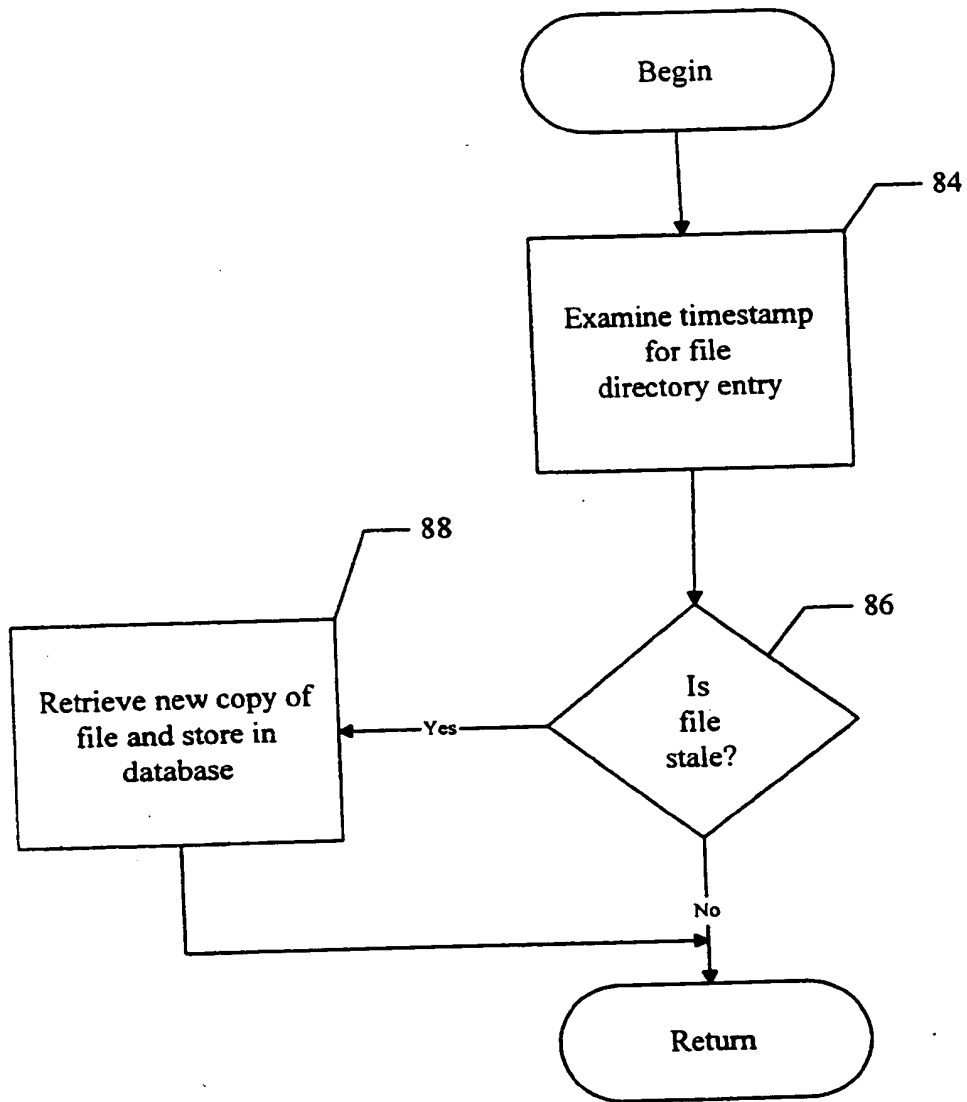


FIG. 7

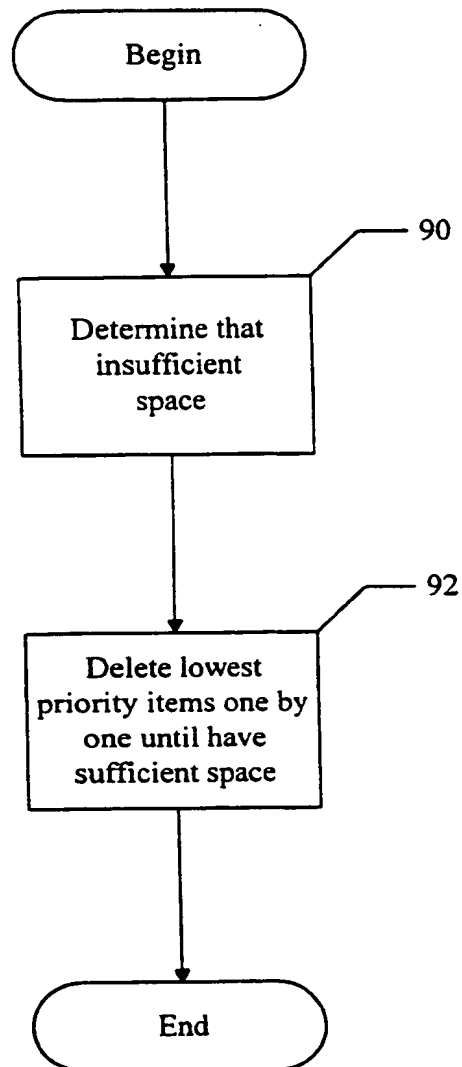
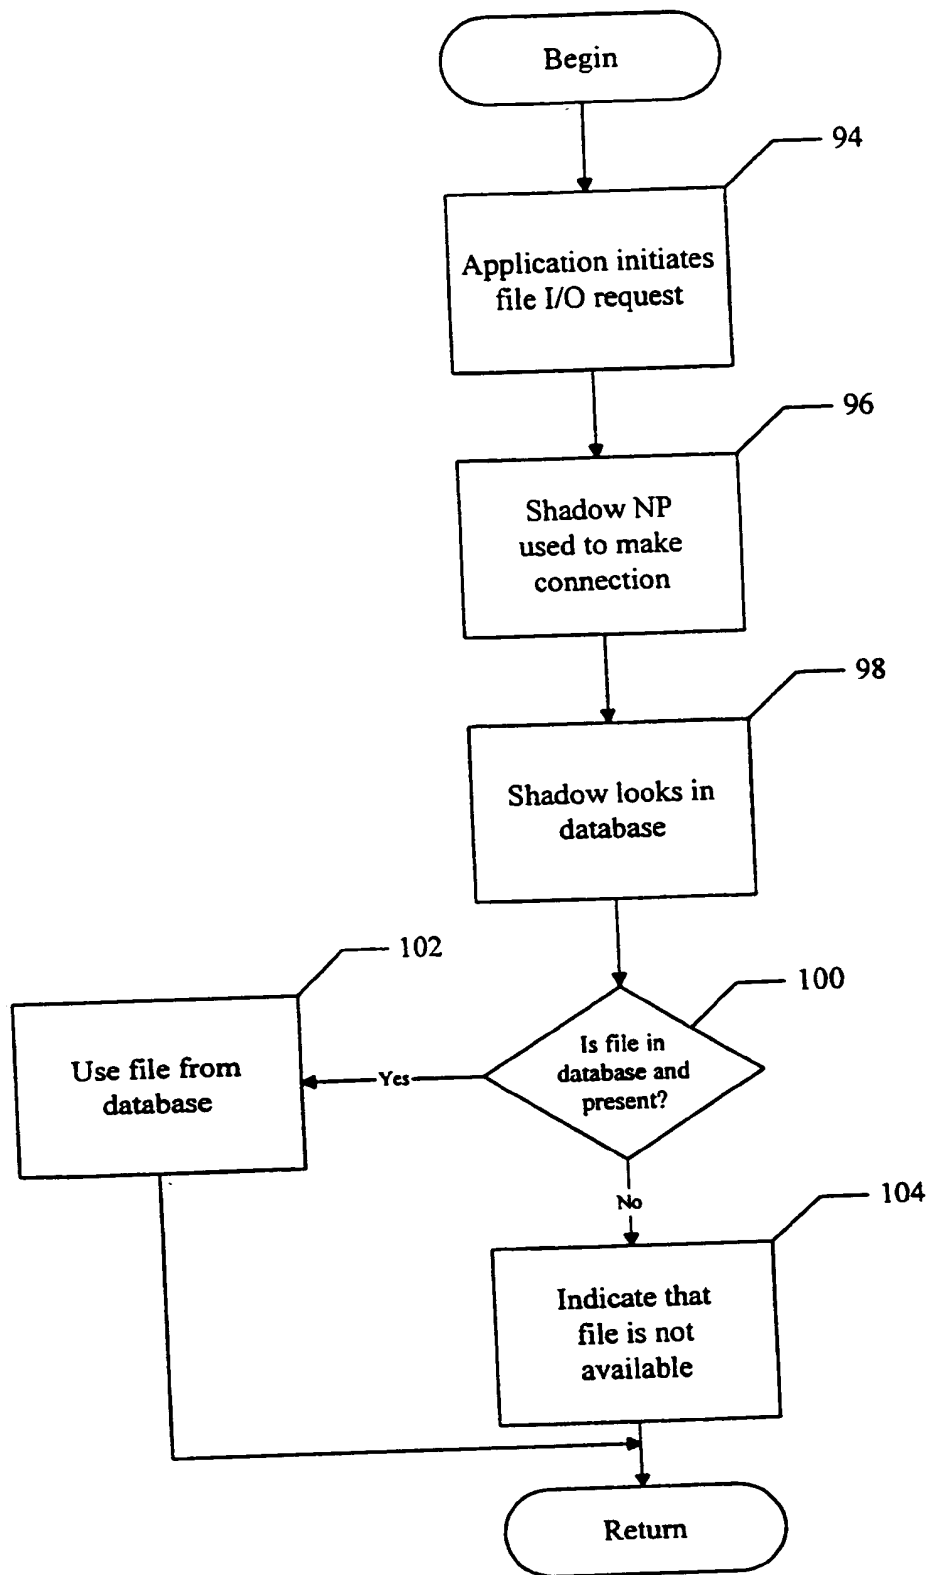



FIG. 8

**FIG. 9**



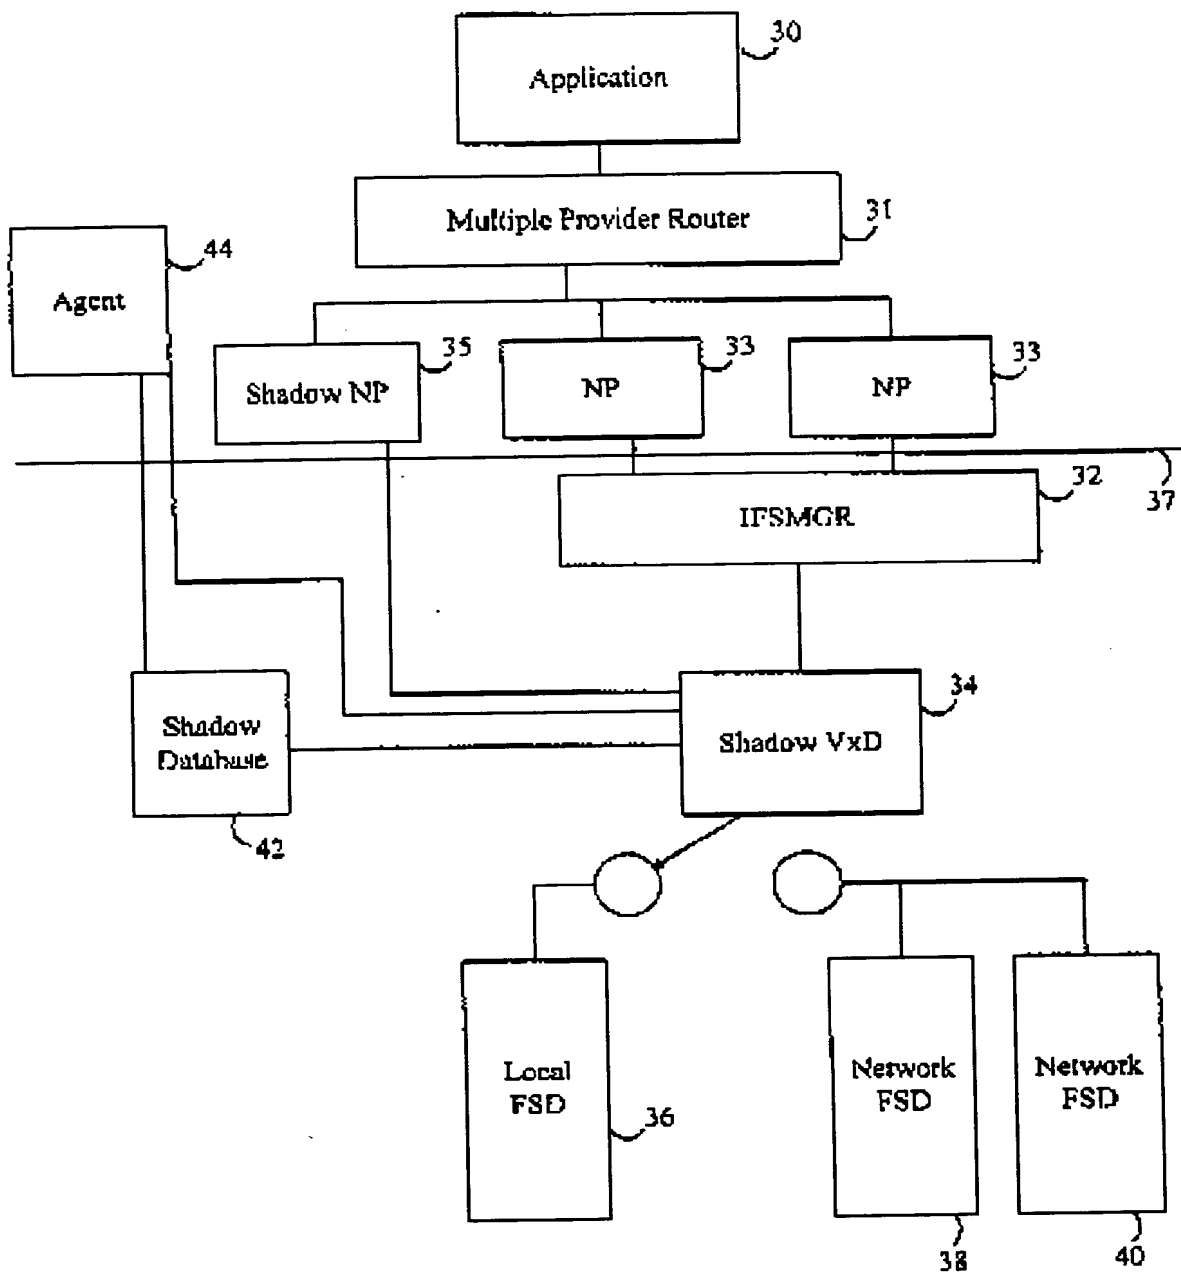
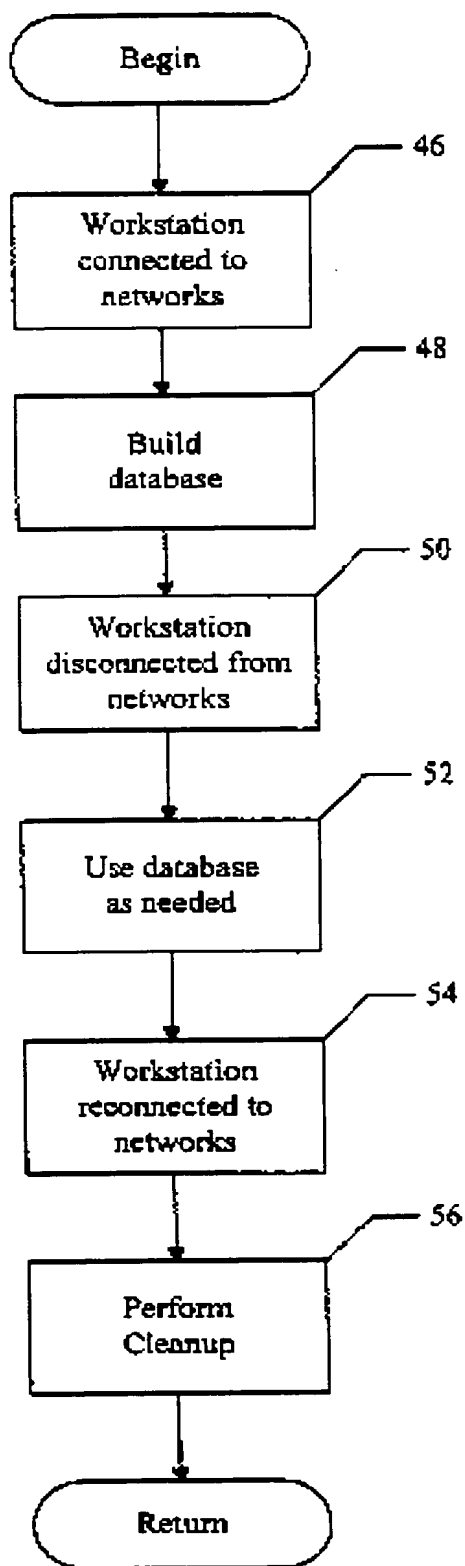


FIG. 2

**FIG. 3**

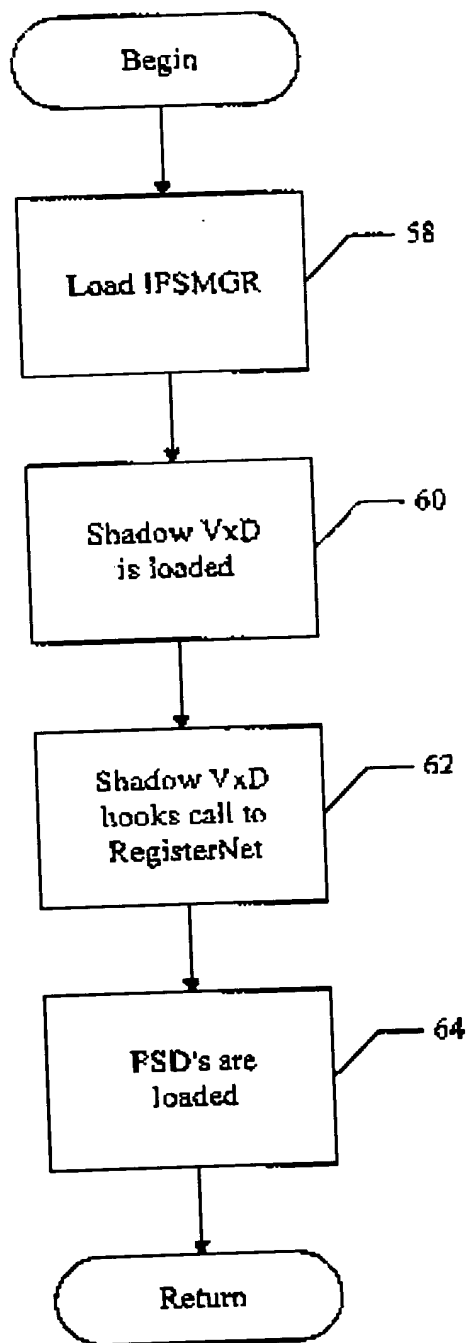

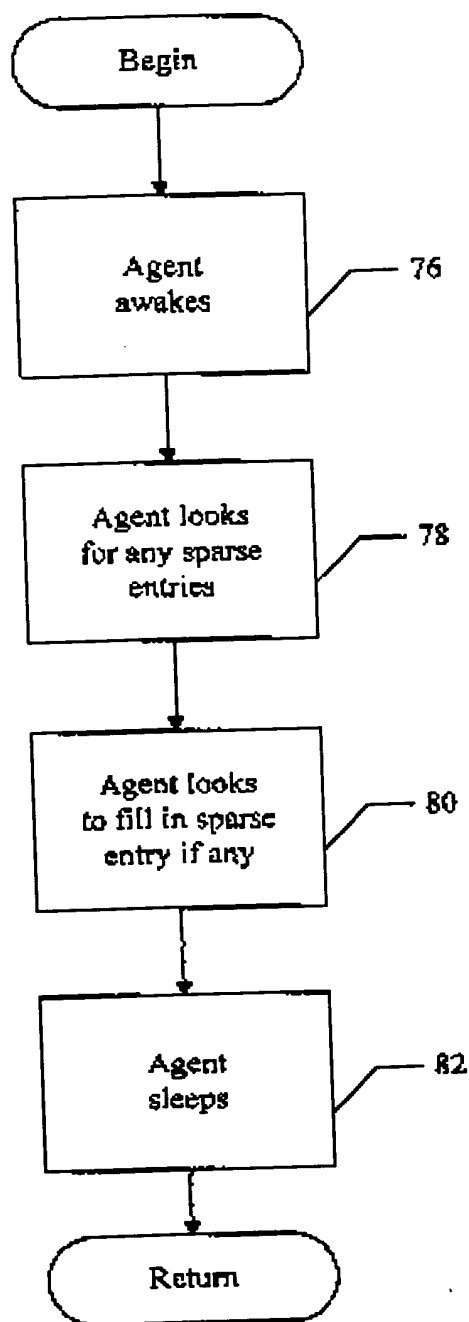
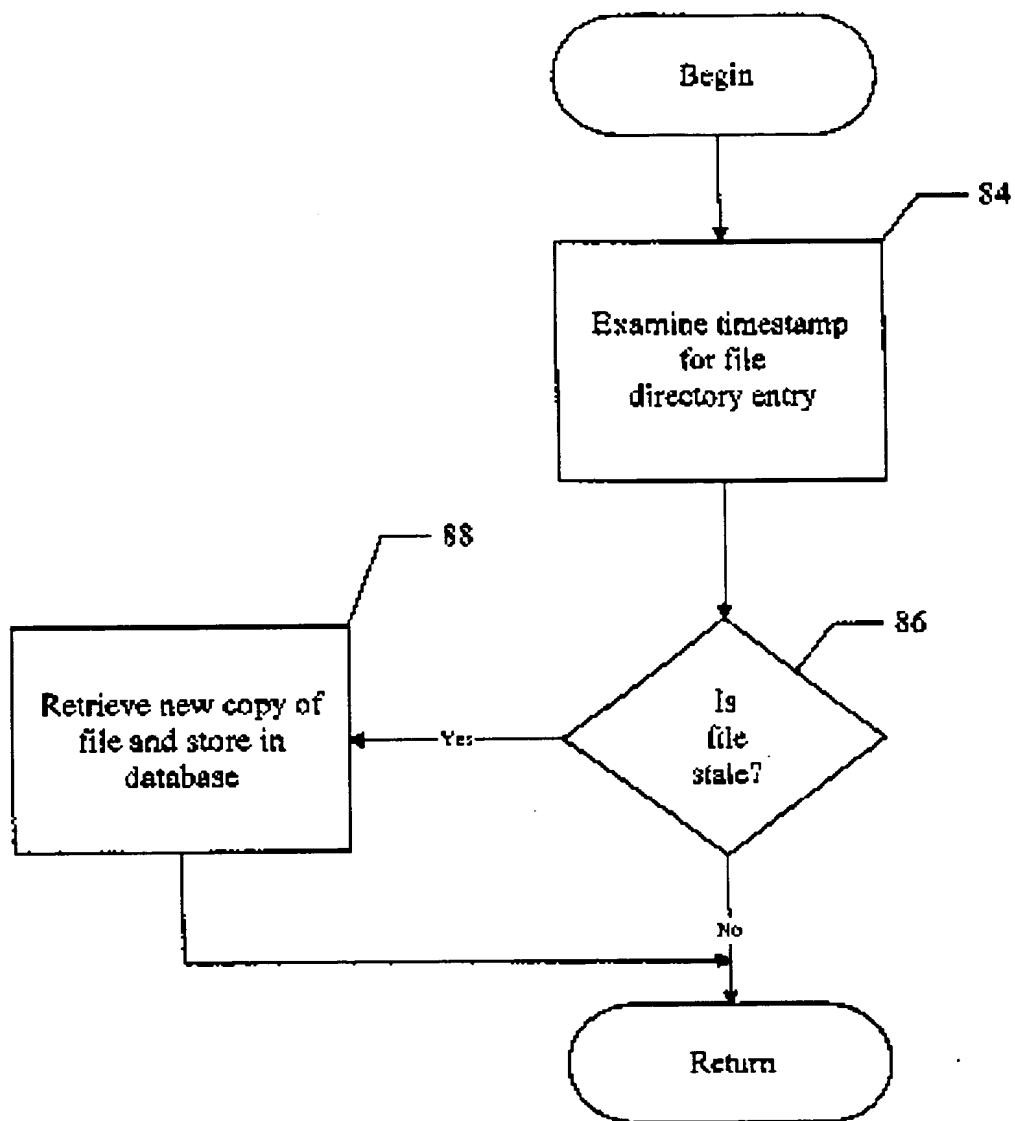


FIG. 4

**FIG. 5**

**FIG. 6**

**FIG. 7**

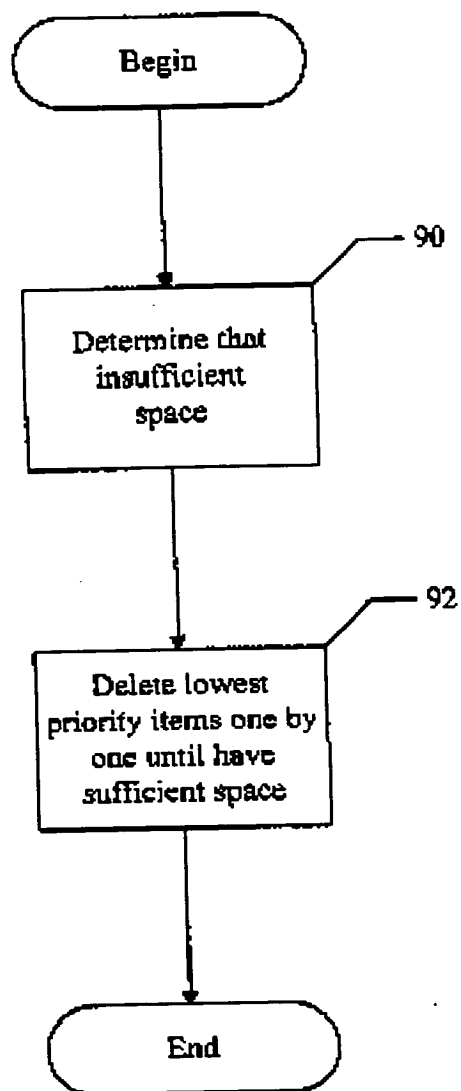




FIG. 8

**FIG. 9**

THIS PAGE BLANK (USPTO)



(19)

Europäisches Patentamt
European Patent Office
Office européen des brevets



(11)

EP 0 713 183 A3

(12)

EUROPEAN PATENT APPLICATION

(88) Date of publication A3:
02.10.1996 Bulletin 1996/40

(51) Int. Cl.⁶: **G06F 17/30**, **G06F 12/08**

(43) Date of publication A2:
22.05.1996 Bulletin 1996/21

(21) Application number: **95117949.8**

(22) Date of filing: **14.11.1995**

(84) Designated Contracting States:
DE FR GB

(30) Priority: **18.11.1994 US 342127**

(71) Applicant: **MICROSOFT CORPORATION**
Redmond, Washington 98052-6399 (US)


(72) Inventor: **Pardikar, Shishir**
Redmond, Washington 98052 (US)

(74) Representative: **Grünecker, Kinkeldey,**
Stockmair & Schwanhäusser
Anwaltssozietät
Maximilianstrasse 58
80538 München (DE)

(54) Network independent file shadowing

(57) Network independent file shadowing is provided by the present invention. The file shadowing mechanism automatically and transparently stores shadow copies of remote file system structures when they are accessed by a computer. The shadow copies of the file system structures are stored within a shadow

database that resides within local memory of the computer. When the computer becomes disconnected from a network, shadow copies of file system structures for the network are used to service requests to access such file system structures.



EP 0 713 183 A3



European Patent
Office

EUROPEAN SEARCH REPORT

Application Number
EP 95 11 7949

DOCUMENTS CONSIDERED TO BE RELEVANT			Relevant to claim	CLASSIFICATION OF THE APPLICATION (Int.Cl.6)
Category	Citation of document with indication, where appropriate, of relevant passages			
X	ACM TRANSACTIONS ON COMPUTER SYSTEMS, vol. 10, no. 1, 1 February 1992, pages 3-25, XP000323223 KISTLER J J ET AL: "DISCONNECTED OPERATION IN THE CODA FILE SYSTEM" * page 11, paragraph 4.3 - page 16, paragraph 4.5 *	1-12, 14-17	G06F17/30 G06F12/08	
A	IEEE TRANSACTIONS ON SOFTWARE ENGINEERING, JUNE 1993, USA, vol. 19, no. 6, ISSN 0098-5589, pages 613-624, XP000394147 RAO H C ET AL: "Accessing files in an Internet: the Jade file system" * page 617, left-hand column, paragraph C - paragraph III *	1-17		
A	COMPUTER, vol. 23, no. 5, 1 May 1990, pages 9-18, 20/21, XP000128601 SATYANARAYANAN M: "SCALABLE, SECURE, AND HIGHLY AVAILABLE DISTRIBUTED FILE ACCESS" * page 17, middle column, line 38 - page 20, left-hand column, line 18 *	1-17		
A	WORKSHOP ON WORKSTATION OPERATING SYSTEMS OCTOBER 14-15, 1993, NA PA, CA, US, vol. 4, 1 January 1993, pages 61-65, XP000505438 NELSON M N ET AL: "GENERIC SUPPORT FOR CACHING AND DISCONNECTED OPERATION" * the whole document *			
The present search report has been drawn up for all claims				
Place of search THE HAGUE		Date of completion of the search 31 July 1996	Examiner Fournier, C	
CATEGORY OF CITED DOCUMENTS X : particularly relevant if taken alone Y : particularly relevant if combined with another document of the same category A : technological background O : non-written disclosure P : intermediate document			T : theory or principle underlying the invention E : earlier patent document, but published on, or after the filing date D : document cited in the application L : document cited for other reasons A : member of the same patent family, corresponding document	

EPO FORM 150 (3.92) (P4031)

**This Page is Inserted by IFW Indexing and Scanning
Operations and is not part of the Official Record**

BEST AVAILABLE IMAGES

Defective images within this document are accurate representations of the original documents submitted by the applicant.

Defects in the images include but are not limited to the items checked:

☐ BLACK BORDERS

☒ IMAGE CUT OFF AT TOP, BOTTOM OR SIDES

☐ FADED TEXT OR DRAWING

☐ BLURRED OR ILLEGIBLE TEXT OR DRAWING

☐ SKEWED/SLANTED IMAGES

☐ COLOR OR BLACK AND WHITE PHOTOGRAPHS

☐ GRAY SCALE DOCUMENTS

☐ LINES OR MARKS ON ORIGINAL DOCUMENT

☒ REFERENCE(S) OR EXHIBIT(S) SUBMITTED ARE POOR QUALITY

☐ OTHER: _____

IMAGES ARE BEST AVAILABLE COPY.

As rescanning these documents will not correct the image problems checked, please do not report these problems to the IFW Image Problem Mailbox.

THIS PAGE BLANK (USPTO)